

Terrain Modelling for In-situ Activity Planning and Rehearsal for the Mars Exploration Rovers

John Wright, Ashitey Trebi-Ollennu, Frank Hartman,
Brian Cooper, Scott Maxwell, Jeng Yen, Jack Morrison
Jet Propulsion Laboratory(JPL)
California Institute of Technology
M/S 264-522 4800 Oak Grove Drive
Pasadena, CA USA 91109-8099

Abstract *Immersive environments are being used to support mission operations at the Jet Propulsion Laboratory. This technology contributed to the Mars Pathfinder Mission in planning sorties for the Sojourner rover and is being used for the Mars Exploration Rover (MER) missions. The stereo imagery captured by the rovers is used to create 3D terrain models, which can be viewed from any angle, to provide a powerful and information rich immersive visualization experience. These technologies contributed heavily to both the mission success and the phenomenal level of public outreach achieved by Mars Pathfinder and MER. This paper will review the utilization of terrain modelling for immersive environments in support of MER..*

Keywords: Immersive environments, MER, terrain modelling, visualization, Mars.

1 Introduction

Current developments in immersive environments for mission planning include several tools utilizing terrain visualization, which make up a system for performing and rehearsing missions. This system, known as the Rover Sequencing and Visualization Program (RSVP)[1][2], includes tools for planning long range sorties for highly autonomous rovers, tools for planning operations with robotic arms, and advanced tools for visualizing telemetry from remote spacecraft and landers. In addition, a Web-based tool, known as the Science Activity Planner (SAP), allows for collaboration by remote scientists in designating features of interest in a similarly immersive environment, using similar terrain models.

As rovers can range over greater distances with more autonomy when they have accurate self-locating systems, the operator paradigm shifts from a hands-on micromanagement level to a hands-off level of mission specification. This calls for a more immersive interaction with the environment with tools for designating waypoints, samples to be collected, regions of hazard and interest, and other types of features. This type of environment is applicable to both rover navigation and operations with

robotic arms and sensors. This type of immersive environment is critical for maximizing operator understanding of the environment for reducing risk and optimizing science return[3][4] and is extremely dependent on terrain models of high quality.

1.1 Requirements

The primary goal of the immersive mission planning tools in RSVP is to provide the mission planners with the best possible understanding of the region. This requires building three-dimensional terrain models from the stereo imagery captured by the three primary camera systems on the rovers, the PanCam, the NavCam, and the front and rear HazCams. Processing of the data from these different camera systems produces three-dimensional models of different resolution and spectral characteristics. Some models may be used independently but it is necessary to combine the models into a single, coregistered dataset to support free-roaming visualization and exploration of the mission area. A further requirement is to support multiple resolutions in the final model products for increased rendering performance, as well as producing models in a variety of formats to be utilized by different applications. Thus, there are requirements to produce individual terrain "wedges" from stereo pairs, register the wedges to each other, merge the wedges into a coherent model, and then produce terrain models tailored for different immersive visualization applications.

In addition, there are specific metrics that must be met in order to adequately support operations. The first is to be able to identify the location of terrain features located within 20m of the rover to within +/-15% of the range from the rover. The second is to identify the location of terrain features located within 20 to 100m from the rover to within +/-25% of their range.

The process descriptions in this paper outline the production of terrain models to support immersing the operator in the environment of another planet, body, or

space to make the mission planning function more intuitive and effective.

2 Building Terrain Models

The System for Unifying Multiresolution Models and Integrating Three-dimensional Terrains (SUMMITT) task has the goal of developing the underlying modelling technology for supporting missions involving rovers. Three-dimensional models of terrain areas are an invaluable asset in planning operations and in reviewing the predicted and telemetered operations of a robot arm. The SUMMITT task had the initial goal of supporting the Mars Volatiles And Climate Surveyor (MVACS) team during the Mars '98 mission, which unfortunately failed. The next Mars surface operations missions were the Mars Exploration Rover missions that landed two rovers. These missions expected to have orbital imagery from Mars Global Surveyor (MGS) and Mars Odyssey, descent imagery from the lander, and local imagery from four sets of stereo imagers on each rover. These sources of imagery were to be combined to create a multiresolution terrain model with very high resolution detail available within the immediate area of operations of the rover. Data constraints, including the limited amount of descent imagery and the wide range of exploration of the rovers, has limited the use of the modelling technology to the data collected by the rovers themselves. Despite this, the mission has had spectacular success.

Figure 1 illustrates the processing performed on the stereo pairs from the PanCams, NavCams, and HazCams. The PanCams and NavCams are mounted on the mast, about 1.5m above the terrain. The HazCams are mounted below the rover deck, about 0.5m above the terrain. Each has a left and right imager generating stereo pairs for evaluating range information and producing three-dimensional models. Each stereo pair is first processed with a correlator that produces a disparity map that identifies matching features in each image. The camera model is then used to compute the range to each pixel in one image using the disparity to its matching pixel in the other image. Then the camera pointing information is used to project the pixels to an (x,y,z) location in three-dimensional space. These XYZ images contain three bands with the x coordinates of each point sample stored in one band, the y coordinates in a second band, and the z coordinates in a third band.

The XYZ images can be thought of as a cloud of point samples generated from a single image pair. These point samples can then be converted to a simple polygon model by merely connecting adjacent points as determined by adjacency within the XYZ image. Such a process creates a wedge of terrain model with the narrow end of the wedge pointing toward the rover. The SAP tool utilizes a set of wedges as a model and manages and renders the individual wedges appropriately for visualization by the scientists. However, for higher performance rendering, a unified

terrain model, with multiresolution characteristics, is desired. The terrain models are registered and merged into an octree. The registration process is very simplistic in that

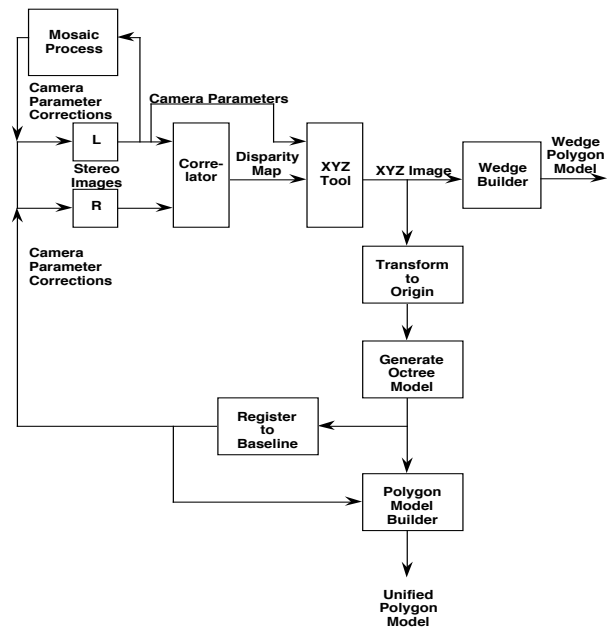


Figure 1 – Terrain Modelling Process Flow

the camera pointing information is assumed to be correct and thus the (x,y,z) coordinates of the samples is accurate. During Mars Pathfinder, such was not the case. However, for MER the pointing is much more accurate and the modelling process is helped greatly.

The imagery from the NavCams, the HazCams, and the PanCams is partially processed independently, then combined with data derived from the other imagery to create the multiresolution terrain models. The registration of different terrain model pieces generated independently from different data sources is a fundamental problem in graphics and modelling. The method explored at JPL uses volumetric primitives (voxels) to represent the terrain to be matched. Voxels have some advantages over the polygonal surface matching methods in that it is easier to represent unknown volumes, such as regions occluded by rocks and hills, and easy to use, multiresolution data structures are available in which to combine the models once they are matched. The registration of the voxel datasets uses an iterative closest points method based on the work of [5]. The entire voxel model is stored in an octree structure, which supports multiresolution data and rapid access while utilizing significantly less memory than a three-dimensional grid. For the MER missions, only imagery from the rovers themselves was available which greatly simplified the registration process. Unfortunately, this process, while widely used in the literature, did not perform adequately to support the MER mission and more work is needed in this area. The fallback method was to utilize the camera pointing with no corrections to define the registration parameters.

The point sample nature of the original data is modified to fit the octree concept which is volumetric in nature. Each point sample is given a volume determined by the pixel field of view of the associated instrument and the range to the sample. This volume is associated with the sample and used to locate the appropriate level of the octree. As data from multiple XYZ images is added, the octree absorbs them and results in a single model of the terrain in the region.

2.1 Model Products

Due to the ubiquity and performance of polygon rendering systems, both software and hardware, a polygon model format was desired. The primary requirements of the polygon model are high fidelity in the terrain imagery with high rendering performance. Because the operator will be making planning decisions that require detailed local knowledge combined with general understanding of more distant terrain, the polygon models must be multiresolution also, or at least reflect the multiresolution nature of the underlying data samples. The use of rendering algorithms and tools in the original point sample space has been considered using such tools as the Volumizer from Silicon Graphics or Mitsubishi Electric's VolumePro real-time hardware. Unfortunately, these systems are optimized for medical data, the primary volumetric data source today. Thus, the datasets are three-dimensional grids of limited size (typically 512x512x512). The datasets expected on the Mars missions will be equivalent to 512^3 for each dataset generated from a single stereo pair. The baseline terrain may be 1024x1024 with tens or hundreds of higher resolution inserts. To avoid the memory usage required by a full 3D grid, an octree was utilized for intermediate storage of the point samples. Current visualization tools do not work with specialized data structures such as an octree. Future work in exploring the use of such techniques as splatting [6] may prove effective.


The multiresolution nature of the point samples precludes the use of a simple algorithm such as Marching Cubes [7] for converting to a polygon surface. Other, more sophisticated methods such as Marching Triangles [8] also do not work well in this environment. To extract the polygon model from the octree, a unique combination of image space sampling and octree space sampling is used. It is desired to produce a tiled, multi-resolution polygon model, yet it is difficult to extract a valid surface from the merged point cloud within the octree without dramatically reducing the resolution and removing features that could be critical to safe operations. To accomplish the polygon model extraction, the connectivity of the samples is extracted from the original XYZ images. Each XYZ image produces one or more independent polygon mesh sections. The images are downsampled to produce lower resolution versions of each section. To achieve tiling, however, the octree is used to divide the points within each XYZ image into separate square regions in X and Y and, structurally,

each section of an XYZ image, at all resolutions, is associated with a tile for that region. To avoid artifacts along the edges of the tiles, polygon edges that cross tile boundaries have both endpoints in both tiles, thus duplicating the polygons along the edges. While not optimal, the increased polygon count is minimal and does not affect rendering performance.

An additional product is the height map. The (x,y,z) values of the point samples is used to populate an image where a z value is stored at each pixel of the image and the (x,y) values are used to index the row and column of the associated pixel. Holes due to missing data and mismatches in resolution and alignment are filled in with a variety of algorithms to produce a smooth product without modifying any of the measured data.

2.2 Terrain Model Usage

Each type of imager on the rovers is used to produce terrain models for a specific purpose. The HazCams are used to produce models for planning and rehearsing operations of the robot arm, the Instrument Deployment Device or IDD. Because the HazCams are on a fixed mount below the deck, they offer the best view of the IDD work volume and are not subject to any pointing errors that might result in the misplacement of an IDD target. Figures 2 and 3 are views of a HazCam model being used to rehearse IDD operations. The NavCam imagery is used primarily for producing terrain models for planning traverses. The NavCams have about a 45° field of view so several of these are combined to model the terrain around the rover. An entire panorama of NavCam models constitutes a site within which several days worth of investigations may be performed. Figure 4 shows a portion of the complete panorama model from the site produced right after driving off the lander. The PanCams are typically used for science observations but are also used to produce terrain models at






Figure 4 – NavCam Mesh

a distance along the expected direction of the next traverse. These higher-resolution images, with a field of view of about 16° , produce good quality meshes out to a greater distance providing for longer traverses being planned safely. Figure 5 shows a merged dataset with both PanCam and NavCam terrain data. Note that where the NavCam data begins to lose quality, the PanCam data is still very good and useful for planning traverses.

2.3 MER Operations

The standard paradigm for MER operations is to drive some distance, declare a new site, capture a panorama of imagery, build terrain models, and then spend some number of days exploring the site. All coordinate information is specified in site coordinates, the origin of which is declared to be the rover's location when the site was declared. As the rover wanders, new sites are declared and panoramas captured. Figure 6 shows how the sites are declared. By Sol 400, 105 sites had been declared for the Spirit rover. Figure 7 shows a view of the traverses performed up through site 17 for the Spirit rover, created




Figure 6 – Successive Site Declarations

by Ron Li, Ohio State University.

3 Model Validation

The terrain models produced for MER are used both individually and in a merged form. Therefore, the validation methodology must reflect this and prove the suitability of the models for short-range arm operations and instrument placement, mid-range navigation and obstacle avoidance, and long-range path planning.




Figure 5 – Merged PanCam and NavCam mesh

3.1 Ground Tests

The validation process began with data collected during one of the Operational Readiness Tests (identified as PORT4-5 in the accompanying tables and figures). This data was collected by identifying a set of rocks in the test area, surveying their locations, and comparing the surveyed locations to the locations of the rocks within the terrain model products. All of the tools that measure terrain locations, surface normal, and other aspects are based on the height map described above rather than the mesh products. The meshes are used primarily for visualization and for arm and instrument collision detection with the

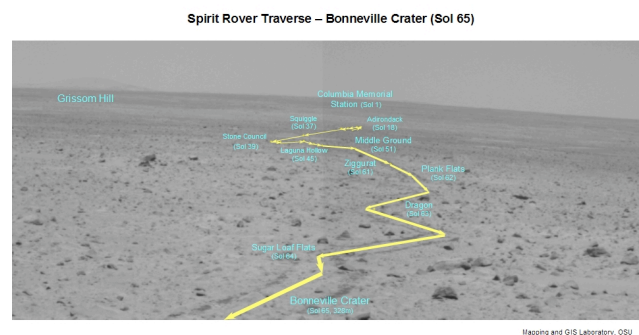


Figure 7 – Traverse History

terrain.

Initially the data indicated a very poor quality. Analysis of the results identified three likely sources of error: calibration of the camera models, calibration of the camera pointing encoders, and insufficiently precise specification of the comparison points. Additional calibration of the camera models and pointing encoders was performed and

Table 1 – PORT6 Model Error Summary in meters

	Height Error	Cross-Range Error	Down-Range Error
Mean	0.022	0.115	0.182
Sigma	0.0096	0.054	0.126
3*Sigma	0.0295	0.344	0.545
Mean+3*Sigma	0.052	0.458	0.727

new data was collected during the next Operational Readiness Test (identified as PORT6 in the accompanying tables and figures). In addition, specific points on the rock targets were selected for surveying and this information used to match to the same points within the terrain models. These results are listed in Table 1 and Figure 8 shows a plot of downrange error versus range to the sample point for both datasets. These indicate that the quality of the models is greatly improved and exceeds our requirements of no more error than +/-15% of range to the sample point in the near to mid-range of 0-20m. From the data collected and reviewed, it was shown analytically that the error in the far-range of 20-100m would be less than 22%, thus also meeting our requirements for performance.

3.2 Performance in Operations

Operational performance on the Martian surface has been

verified in several ways. For near-range data captured by the HazCams and used for defining and rehearsing arm and instrument placement, a comparison of expected and actual instrument contact locations was performed. Because the arm has commandability and repeatability of less than 1.0mm, any differences in contact locations are primarily due to camera model issues, particularly if the ground is level and slippage is minimal. It was found that some camera model issues were present on Spirit while Opportunity was performing very well. The source of the problems may have been jarring during traverses or flight, or contractions of the material due to the cold of the Martian surface and the large temperature swings. The




Figure 9 – Simulated vs Measured Pitch and Roll




Figure 8 – Range Errors for Two Tests

camera models were recalibrated by placing the arm instruments at known positions (errors less than 1mm in arm positioning), capturing a series of images, and computing new camera models. This process produced new camera models that have been used for processing the HazCam data with smaller error. All instrument placement commanding is done with 1cm of overdrive to allow for errors in terrain modelling and has never failed to make contact within the 1cm tolerance. Thus, it has been shown that the range error is less than 1% in the HazCam models where the range is approximately 1m.

An additional field check of the terrain model quality was performed by comparing the attitude of the simulated rover driving over the terrain model to that of the actual rover as it drove over the Martian terrain. Figure 9 shows the comparison of the measured roll and pitch reported by the rover as it drove and the roll and pitch measured during simulation. Since the simulation does not adequately model slippage nor does it model decisions made by the onboard hazard avoidance, the rover's XY position is determined by localization. This process utilizes imagery captured by the rover to triangulate on features to determine the rover's position. Then the position and heading information is fed into the simulation software to determine pitch and roll based on kinematic settling on the modelled surface. While the curves do not align exactly, there is good correspondence between simulated and measured values. The misalignment can be due to other flight software activities varying the timing of the mobility activities, and other factors.

4 Conclusions

The Jet Propulsion Lab is utilizing enhanced immersive technologies for supporting MER mission operations. Immersive technologies and systems are aiding the operations teams in making mission critical decisions. Creating models of the operational environment and providing visualization tools to explore and interact with that environment are the key aspects of the MER mission tool suite. have made MER so successful.

The processes described here have been shown to produce high quality terrain models with the resolution and accuracy necessary to support mission operations on the Martian surface. Activity commanding based on the terrain knowledge gained from interacting with these models has been very successful.

Immersive technologies will continue to offer more capability to mission operations teams in the foreseeable future. This will be especially important as exploratory craft gain in autonomy and intelligence and begin to explore farther afield from their initial landing site. Operations will change from precisely specifying number of revolutions of each wheel to broadly defining goals and strategies. Such paradigms require a broader understanding

of the environment to facilitate rapid decision making and terrain models are key components in this process.

Acknowledgements

The work described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

- [1] Wright, J., Hartman, F., and Cooper, B. (1998) Immersive Environments for Mission Operations: Beyond Mars Pathfinder, Proceedings of SpaceOps '98, Tokyo, Japan, 1998.
- [2] Cooper, B. (1998) Driving on the Surface of Mars Using the Rover Control Workstation, Proceedings of SpaceOps '98, Tokyo, Japan, June, 1998.
- [3] Ellis S.R. (1994). What are Virtual Environments? IEEE Computer Graphics and Applications. 17-22.
- [4] McKenna, M. & Zeltzer, D. (1992). Three Dimensional Visual Display Systems For Virtual Environments. Presence: Teleoperators and Virtual Environments. 1(4). Cambridge, MA: MIT Press. 421-458.
- [5] Champeboux, G., Lavallee, S., Szeliski, R., and Brunie, L., (1992) From Accurate Range Imaging Sensor Calibration to Accurate Model-Based 3-D Object Localization, Proceedings of IEEE CVPR '92, pp. 83-89, 1992.
- [6] Swan, J.E II, Mueller, K., Moller, T., Shareef, N, Crawfis, R., and Yagel, R. (1997). An Anti-aliasing Technique for Splatting. IEEE Visualization '97. Phoenix, AZ: IEEE. 197-204
- [7] Lorensen, W. and Cline, H. (1987) Marching Cubes: A High Resolution 3D Surface Construction Algorithm. Proceedings of SIGGRAPH '87, pp. 163-169, July 1987.
- [8] Hilton, A., Stoddart, A.J., Illingworth, J., and Windeatt, T. (1996) Marching Triangles: Range Image Fusion for Complex Object Modelling. Proceedings of International Conference on Image Processing, 1996.